LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

The Hyperion Project: Partnership for an Advaned Technology Cluster Testbed

M. Seager, M. Leininger

April 28, 2008

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

The Hyperion Project:
Partnership for an Advanced Technology Cluster Testbed

Matt Leininger[†], Mark Seager, Brent Gorda
Lawrence Livermore National Laboratory[1]
LLNL-TR-403271

Partners from:

Cisco, Data Direct Networks, Dell, Intel, LSI, Mellanox, Qlogic, RedHat, SuperMicro, Sun

Version 11
April 28, 2008

1. Executive Summary

The Hyperion project offers a unique opportunity to participate in a community-driven testing and development resource at a scale beyond what can be accomplished by one entity alone. Hyperion is a new strategic technology partnership intended to support the member-driven development and testing at scale. This partnership will allow commodity clusters to scale up to meet the growing demands of customers multi-core petascale simulation environments. Hyperion will tightly couple together the outstanding research and development capabilities of Lawrence Livermore National Laboratory with leading technology companies, including Cisco, Data Direct Networks, Dell, Intel, LSI, Mellanox, Qlogic, RedHat, SuperMicro and Sun. The end goal of this project is to revolutionize cluster computing in fundamental ways by providing the critical software and hardware components for a highly scalable simulation environment. This environment will include support for high performance networking, parallel file systems, operating system, and cluster management. This goal will be achieved by building a scalable technology cluster testbed that will be fully dedicated to the the partners and provide:

---

1. A scalable development testing and benchmarking environment for critical enabling Linux cluster technologies
2. An evaluation testbed for new hardware and software technologies
3. A vehicle for forming long term collaborations

2. Background

Over the past several decades, parallel computing has emerged as a contemporary, and indispensable tool in science and engineering. Current state-of-the-art simulation environments have enabled key break-throughs in materials science, biology, chemistry, and physics. Lawrence Livermore National Laboratory (LLNL) is a leader in designing and deploying commodity Linux clusters to successfully deliver economic capacity computing cycles to our NNSA, DOE, and institutional programs. The typical LLNL capacity computing simulation environments consists of commodity processors, memory, and motherboard hardware connected through a high-speed low-latency InfiniBand interconnect. Within the DOE NNSA laboratories these large-scale clusters have grown to be 1,152 nodes with 18,432 cores (LLNL Juno) and 4,480 nodes with 8,960 cores (SNL Thunderbird). Each of the NNSA Tri-labs is currently managing over 6,000 Linux cluster nodes that are dedicated mission critical production capacity computing resources. The experiences with large-scale clusters have enable the development of next generation simulation environments. Future multi-petascale simulation environments will utilize many of the features found in Hyperion, including multi-core processors, IB/Ethernet SAN, Linux system software, parallel file systems, and high performance storage solutions. Hyperion will have a direct impact on the design of the next generation storage system for the NNSA multi-petascale Sequoia platform.

Large-scale clusters are also having a significant impact on how the commercial and industry sectors deliver computing. Financial services, oil and gas, and other key business sectors have migrated from proprietary solutions to commodity clusters. Clustering has completely transformed the way these organizations value and utilize computation. Although the growth in cluster computing is accelerating, the bottlenecks to improving cluster usability, reliability, and manageability still exist. The current gaps to expanding opportunities for cluster computing within the government, business and, industry sectors include, the prolonged deployment and integration timelines, the stability of system software, and the challenge of managing thousands of cluster nodes.

The deployment of a community accessible scalable cluster environment will lay a strong foundation for eliminating the barriers to future growth in cluster computing. The lack of a scalable cluster development and testing environment dedicated to benchmarking, to software development, to system software stabilization, to scalability testing, and to the evaluation of new hardware and software technologies has been identified by DOE and several commercial organizations as a key missing link to fully realizing the potential of cluster computing. The Hyperion project will deliver this vital scalability testbed environment. The Hyperion cluster will not be used by end application users (i.e. Hyperion *is not* a production cluster for DOE science and engineering applications), but rather will be a generally available platform to the partners for joint R&D efforts.

3. Hyperion Objectives and Goals

The Hyperion project is focused on the goal of providing a scalable technology cluster testbed that will be fully dedicated to the the partners and provide:

1. A scalable development environment for benchmarking and testing of critical enabling Linux cluster technologies
2. An evaluation testbed for new hardware and software technologies
3. A vehicle for forming long term collaborations

The scalable testing environment will be dedicated to software development and stabilization of three key components: the OpenFabrics software stack, parallel file systems, and a scalable clustering stack (NNSA Tripod Open Source Stack - TOSS). The OpenFabrics software stack is the first high performance network stack for Linux that supports the RDMA and low latency technologies for InfiniBand and Ethernet. Parallel file systems have become a essential ingredient for large-scale simulation environments. Lustre is currently widely deployed in production environments at the LLNL, and therefore will be the primary focus for parallel file system work. Emerging storage technologies like parallel NFS (pNFS), GFS, commodity RAID devices, and NAND flash storage devices will also be explored. The NNSA laboratories have recently consolidated our Linux cluster environments to a single software stack called TOSS. TOSS is based on RedHat Enterprise Linux, OpenFabrics, Lustre, cluster management utilities, and a variety of open source and proprietary third party compilers and performance analysis tools. These three main components will provide the Hyperion partnership with a complete software stack that is certified to meet the partnership's performance, scalability, stability, and manageability goals to the scale of at least 1,152 nodes on the Hyperion cluster.

Hyperion will provide an evaluation testbed for new hardware and software technologies. The architecture for Hyperion (described in the next section) provides 576 spare InfiniBand ports for the testing of new hardware, future expansion, technology refreshes, and upgrades. This hardware could include processor, memory, networking, storage, visualization, or other new technologies that partners wish to evaluate. For example, scalable petascale I/O technologies for achieving over 512 GB/s of delivered I/O bandwidth represent one class of new technologies of interests to the Hyperion partners. Scalable networking activities in the areas of MPI collectives, job launch, IB fabric management tools, and new network interconnect technologies for adaptive routing and congestion management are also of interest.

Hyperion enables long-term strategic partnerships between LLNL/NNSA and industry. While providing a development and testing vehicle for the partners, Hyperion also lays a strong foundation for establishing partnerships and collaborations with the OpenFabrics Alliance, RedHat, LLNL Lustre Center of Excellence, the DICE Alliance, other U.S. government agencies, financial services, oil and gas, and other commercial cluster

computing users. Section 7 outlines future extensions and opportunities for collaboration.

4. Hyperion Architecture and Budget

Architecture

The architecture of Hyperion is based on the concept of a scalable unit (SU).   The SU concept was developed by the Tri-labs and successfully implemented in the LLNL Peloton procurement. The SU's are lego-like building blocks that provide a tractable architecture for deploying small, medium, and large-scale Linux clusters.   Single SU's can be deployed as small clusters, a medium system would combine two SU's, and large-scale systems would combine four or more SU's.   LLNL recently deployed four Peloton class platforms including a two SU, a four SU, a six SU, and an eight SU clusters. The largest eight SU cluster is 1,152 nodes, the same size as the proposed Hyperion base system.   The SU concept has greatly improved the usual extended deployment and integration timeline of a large-scale cluster, reducing it from months to days.   The final Peloton cluster, named Minos, is a six SU cluster (864 nodes with 6,912 cores) and was moved into classified production in three weeks after the initial delivery to LLNL. This is an extremely accelerated timeline, relative to past practices.
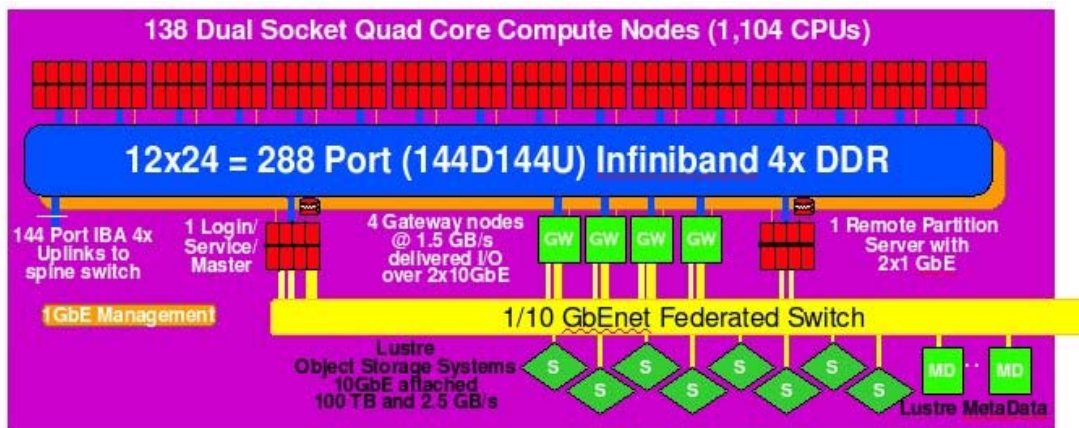

Figure 1: Hyperion Scalable Unit Architecture

The SU architecture is show in figure 1. The SU includes a multi-vendor IB fabric with 288 IB 4x DDR ports that provide a full fat-tree topology with 144 links down to nodes and 144 links up to the next stage of switching. Each SU has 144 nodes with 138 compute nodes, 4 gateway nodes, 1 remote partion server (RPS) node, and 1 login/service/master (LSM) node.   Each node type is a 2-socket quad-core platform with 8 GB of memory.

The gateway nodes route between the InfiniBand cluster fabric and the SAN, which could either be ethernet or InfiniBand. The LSM node connects to the site ethernet

infrastructure and is the interface between the users and the cluster. The RPS node provides the Linux image for the diskless compute and gateway nodes, while also running some system management utilities, and if needed provides swap space to the compute and gateway nodes.   A single SU based on an Intel Harpertown quad-core Xeon 2.66 Ghz CPU chip would provide 12.3 TF/s, 1,152 CPU cores, and 1.15 TB of memory. All nodes utilize PCI-Express generation 2 I/O technology from an Intel Seaburg chipset, thus providing a clear roadmap for upgrading to 4X QDR InfiniBand. It is highly desirable that each node have two PCI-Express generation 2 (PCIe Gen2) x8/x16 slots, or one PCIe Gen2 x8/x16 slot and one InfiniBand HCA down on the motherboard. This will allow for future upgrades of the InfiniBand fabric to dual-plane, one plane 4X DDR and one plane 4X QDR.
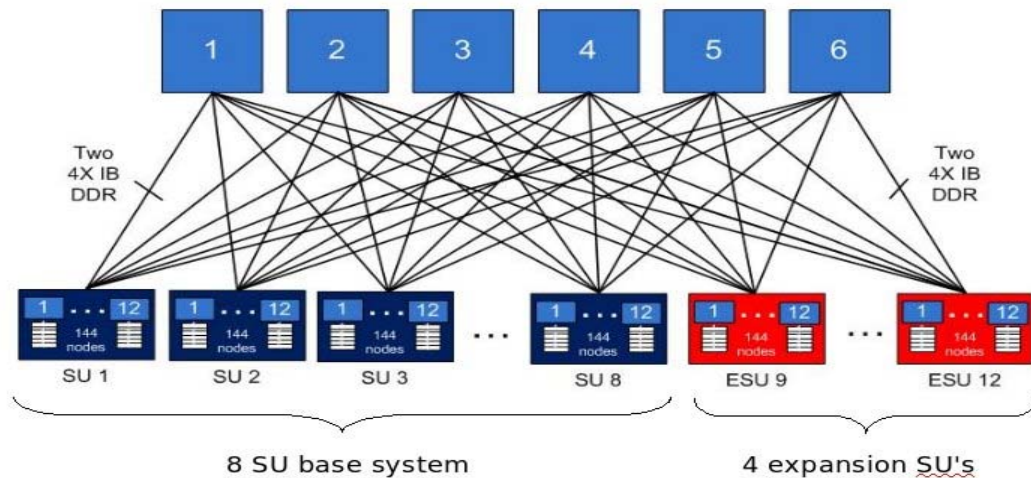

Figure 2: Hyperion InfiniBand Fabric Topology

The Hyperion InfiniBand fabric is a full fat-tree and is designed to support up to 12 SU's , or 1,728 IB ports. The first 8 SU's, or 1,152 IB ports, support the Hyperion base system. The additional 4 SU's (or 576 IB ports) may be used for cluster expansion, new technology evaluation, IB direct attached storage, or any other I/O expansion.   In figure 2, the initial SU's are labeled SU 1-8, and the expansion SU/ports ESU 9-12.

An example SU rack layout is shown in figure 3. The compute nodes are based an standard 1U 2-socket servers, while the gateway, RPS, and LSM nodes are standard 4U 2-socket servers. Higher density solutions are highly desirable.

An example Hyperion storage scalable unit (SSU) consists of two couplets with 480 SATA hard disks. Figure 4 shows an example SSU configuration with four SSU and the corresponding Lustre OSS nodes. This five rack configuration with 750 GB SATA disks would provide 1.152 PB disk and 8-25 GB/s of delivered I/O bandwidth through sixteen Lustre OSS heads/nodes.
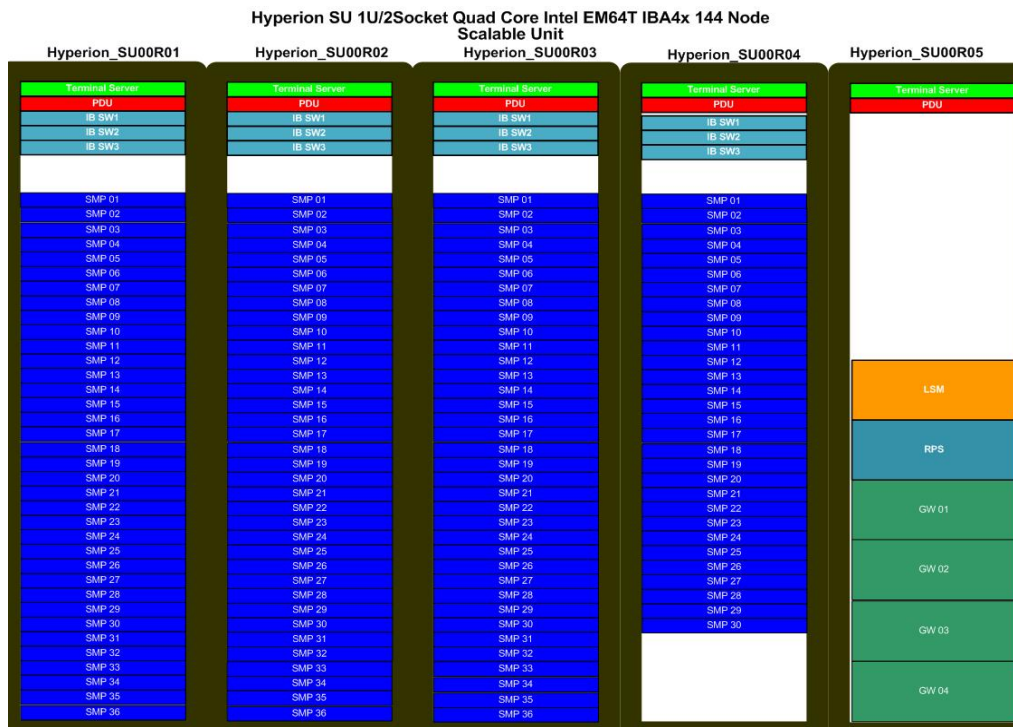
Figure 3: Example Hyperion Scalable Unit Rack Layout

Figure 4: Four Hyperion Storage Scalable Units

Budget

The sustainable plan for the Hyperion budget reflect the investments from both the DOE and industry. Industry partners invest in Hyperion and in return get time on the machine.

The preliminary LLNL cost targets are:
- $3.0 M Base System
  - $2.0 M Nodes, Racks, PDU, InfiniBand HCAs
  - $1.0 M InfiniBand switches and cables, management ethernet
- $750K Storage Scalable Units
- $250K Citing & Installation

Partners Hyperion Responsibilities

1. Initial Investment - The Hyperion industry partners provide funding support to enable the Hyperion project. The estimated total industry investment ranges from $5.0 – $7.0 M depending on the exact configuration of Hyperion.
2. Annual Upgrades - The partners will also contribute to yearly technology refreshes and/or upgrades.
3. Additonal Resources – The partners will provides resources to participate on the operations team and executive board. Partners may also provide technical expertise or subject matter experts as necessary per deployment/project/or maintenance of the system. Resources should be agreed to by all parties up front.

LLNL Hyperion Responsibilities
1. Initial investment – TBD ($3.0-5.1M or possible lease)
2. Annual Contribution: It is estimated that DOE will contribute $500K/yr to technology refreshes/upgrades.
3. Administration of Hyperion (0.75 FTE)
4. Siting of Hyperion
5. Power & Cooling of Hyperion

5. Management Plan

To ensure the goals of the partnership the Hyperion project will provide an execution plan that involves the participation of all partners.   The proposed Hyperion governance model is made up of an operations team and an executive board.   The operation team and executive board are composed of members from each of the Hyperion partners. The operations team is responsible for setting the day-to-day schedule of activities for the use of the Hyperion platform. The operations team has weekly meetings as well as interacting within the community via instant messaging and email. The executive board sets policy, allocations, collaboration interaction models and deals with strategic issues and priorities. The executive board will meet quarterly.

Company	Operations Team Member	Executive Board Member
Cisco	Steven Carter	
Data Direct Networks		
Dell		
Intel		
LLNL	Matt Leininger	Mark Seager
LSI		
Mellanox		
Qlogic		
RedHat		
SuperMicro		
Sun		

6. Benefit to Partners

Hyperion is the link between the public and private sector that paves the way for broad scale adoption of large compute cluster systems. Only through collaboration could all of the interested parties procure and house a system of this magnitude: and only a system of

this magnitude can provide the test bed to develop the technology advances currently missing in both the public and private sectors. There is simply no other way to procure, house, and manage a system of this size without an enormous burden on whoever chooses to do so. The value delivered to the partners in the program will be magnified by the fact that in acting now, it is highly unlikely that any other organization would be able to execute on a plan as ambitious and valuable as Hyperion.

The benefit for partner participation in the Hyperion project can be summed up in four main points:

1. GROW BUSINESS: Demonstrating, to users, the ability to break through current architectural bottlenecks is key to the growth of cluster computing. A scalable test for system software will increase market adoption of cluster technology - The critical barrier to wider market adoption of HPC clusters is the scalability, manageability, and stability of system software. This software gap is hindering the development of advanced HPC application software and limiting the market grown in HPC clusters. The Hyperion project addresses this challenge by forming an ongoing collaboration between LLNL/NNSA and industry leaders in commodity cluster computing, by providing a scalability testing platform to the partners, and enabling a proven software environment that is tested and certified on over 1,000 nodes.

2. REDUCE RISK AND TIME TO MARKET: Scalable, Proven Cluster Management Stack and Stability Among Key Technology Software Components - By delivering a proven software stack that is scalable to a thousand nodes, the Hyperion project will enable the advancement of HPC application software and drive a wider market adoption of HPC commodity clustering technologies to the broad range of commodity cluster sizes. A highly robust OpenFabrics software stack and improved stability of Lustre, pNFS, and GFS will provide a high rate of confidence for this generation and the next generation of scalable clusters. By adopting TOSS, the Hyperion partners get a working scalable cluster distribution with Tri-lab technical involvement. The flexibility of the Hyperion architecture provides a world class hardware and software testing environment that will reduce the time to market of new technologies.

3. CUSTOMER REFERENCE AND BENCHMARKING: The Hyperion project provides a leading edge customer reference site and benchmarking platform for the partners future procurement opportunities. The value is twofold:
   1. Vendors will be better equipped to respond to large procurements requiring benchmarks. Results up to 9,216 cores will be actual results rather than extrapolations, which provides greater opportunity to optimize codes and extract more efficient results.
   2. The system will also dramatically reduce the exposure in meeting acceptance criteria for the system. When benchmarks are actual rather than extrapolations, the degree of accuracy in performance predictions increases dramatically, thereby reducing financial risks associated with acceptance testing.

4. LEADERSHIP ACKNOWLEDGEMENT: Best in Class Partnership Between Government and Industry to Scale - A sustainable partnership between LLNL/NNSA and industry leaders provides a strong coupling and mutual leverage among partners.   Such a collaboration is necessary to provide a focused effort of combined expertise on HPC scaling.   The deployment of a order 1,000 node scalability testbed is a fundamental requirement for advancing HPC system software.   By combining the expertise and sharing costs, the Hyperion partnership can provide this scalability platform that would not otherwise exist at any one organization.